

Singularity Overview

Galen Hunt and James Larus

Microsoft Research

July 17, 2006

MSR Faculty Summit

Large, Diverse Research Team

- Lead by Galen Hunt and Jim Larus
- **MSR Cambridge**
 - Paul Barham, Richard Black, Tim Harris, Rebecca Isaacs, Dushyanth Narayanan
- **MSR Redmond**
 - Advanced Compiler Technology Group:
 - Juan Chen, Qumyran Mangus, Mark Plesko, Bjarne Steensgaard, David Tarditi
 - Foundations of Software Engineering Group:
 - Wolfgang Grieskamp
 - Operating Systems Group:
 - Mark Aiken, Chris Hawblitzel, Orion Hodson, Galen Hunt, Steven Levi
 - Security and Distributed Systems:
 - Dan Simon, Brian Zill
 - Software Design and Implementation Group:
 - John DeTreville, Ben Zorn
 - Software Improvement Group:
 - Manuel Fahndrich, James Larus, Sriram Rajamani, Jakob Rehof
- **MSR Silicon Valley**
 - Martin Abadi, Andrew Birrell, Ulfar Erlingsson, Roy Levin, Nick Murphy, Ted Wobber

"Modern" OS And Applications



"Modern" OS And Applications




- Design parameters
 - scarce resources
 - benign environment
 - knowledgeable and trained users

"Modern" OS And Applications



- Design parameters
 - scarce resources
 - benign environment
 - knowledgeable and trained users



World Changed

- Hardware and software industries were wildly successful
 - machines are fast
 - memory is cheap
 - computers are ubiquitous
- Malicious environment
 - ubiquitous worms, viruses, scams, attacks, ...
- Few users understand computers or software



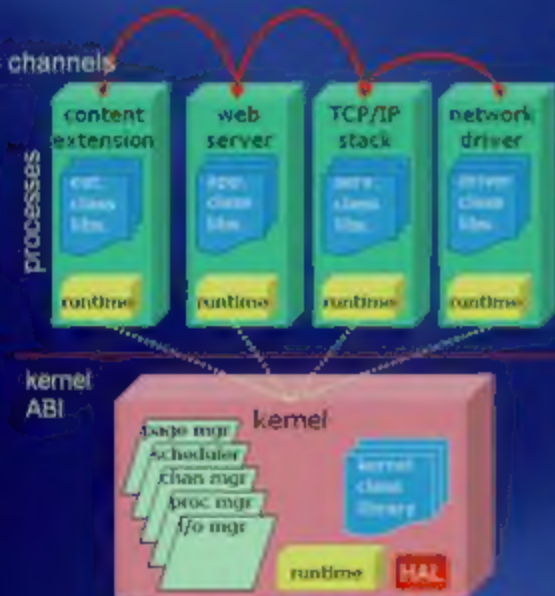
Singularity

- Goal: technology and techniques to build more dependable systems
- Dependable: predictable behavior and easily understood usage model
 - consumer satisfaction: new car vs. new PC
 - car has .99 to .999 availability (9-90 hours down time/yr)
- Research on new OS, languages, and tools
 - attack problem from multiple directions
 - working research prototype (not Windows replacement)
- No magic bullet
 - mutually reinforcing improvements to languages and compilers, systems, and tools

Key Approaches

1. Pervasive use of safe (& analyzable) programming languages
 - type safety and memory safety
 - including device drivers, OS components, applications
2. Improve system resilience despite software errors
 - failure boundaries between components
 - improve extension model
 - explicit error notification
3. Increased verification
 - specification at multiple levels of abstraction
 - closed environments with explicit cross-domain interfaces
 - design for verifiability

Singularity OS



- **Closed Kernel**
 - 95% written in C#
 - 17% of files contain unsafe C#
 - 5% of files contain x86 or C++
 - OS services & drivers in processes
 - kernel closed at boot time
- **Software isolated processes (SIPs)**
 - all user code is verified safe
 - some unsafe code in trusted runtime
 - processes closed at start time
- **Safe and efficient communication via strong interfaces**
 - channels between processes
 - channel behavior is specified & checked
 - checked behavior enables efficient communication
- **Type safety is crux of verification and protection**


Challenge 1:

Pervasive Safe Languages




- Singularity is written in extended C#
 - actually Spec#
(C# + pre/post-conditions and invariants)
- Added features for systems programming
 - increase programmer control over allocation, initialization, and memory layout
- Language design to support programming and verification
 - message passing
 - factoring libraries into composable pieces
 - compile-time reflection

What About The Runtime?



- JVM & CLR's design not always appropriate
 - "one runtime, one size fits all"
 - monolithic, general-purpose environment
 - large memory footprint (~4 MB process for CLR)
 - many dependencies (CLR PAL requires >300 WinSxS APIs)
 - JIT compilation
 - increases runtime size and complexity
 - unpredictable performance
 - replicate OS functionality
 - security, threading, configuration, etc.
 - more is less

Singularity Runtime



Singularity Runtime



Singularity Runtime

Libraries

Application

Singularity
Runtime
(GC, etc.)

Whole Program
Optimization

Application
Library

Singularity Runtime

Libraries


Application

Singularity
Runtime
(GC, etc.)


Whole Program
Optimization



Singularity Runtime



Singularity Runtime



Small, customizable Runtime

Small, customizable environment

- ahead-of-time global optimizing compiler (MIR Barton specializes runtime and libraries)

- eliminate code for unneeded/unused language features, runtime, application/library code

Small, customizable runtime and libraries

- Runtime, garbage collector, and libraries selectable on per-process basis

- e.g. reduce memory and computation overhead

- e.g. enforce design discipline and system policies per process


- Eliminate OS functionality from runtime

- e.g. security, resource allocation, etc.

- Provide OS mechanism for enforcing system policy

- e.g. runtime can constrain behavior (e.g. disallow enforcement)

Singularity Runtime




Runtime Overhead

Memory footprint
"Hello World" process

	Singularity	FreeBSD 5.2	Linux 2.6.11 (Red Hat FC4)	Windows XP (SP2)
C - static lib		232K	554K	544K
C++ - static lib		784K	1,216K	572K
C# - w/ GC	408K*			3,750K

- C# process w/ GC has similar memory footprint to C++
- minimal process (no GC at user level) ~10K

Runtime Resilience



- Software errors should not cause system failure
- Resilient system architecture
 - isolate system components to prevent data corruption
 - provide clear failure notification
 - implement policy for restarting failed component

Process Architectures


- User
- System

Process

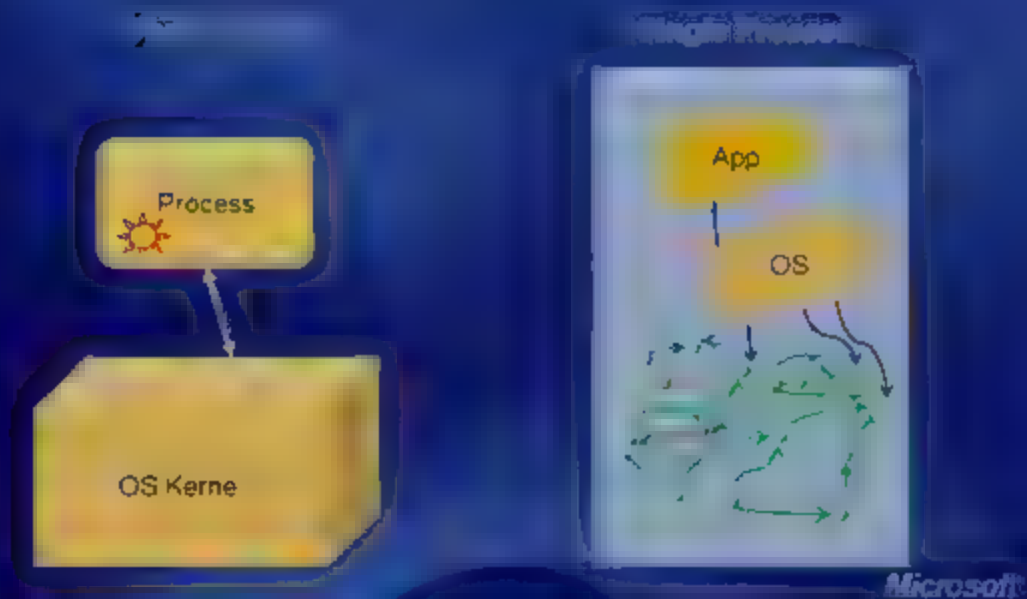
OS Kerne

App


OS




Process Architectures




Process Architectures



Process Architectures



Process Architectures



Open Process Architecture



Open processes

- dynamic code loading and runtime code generation

- DLLs, device drivers, hotfixes, browser plug-ins, device drivers, kernel, etc.

- cross-process memory sharing

- system API allows one process to alter state of another

- Near ubiquitous (Windows, Unix, etc.)

- originated in Multics

- Shared state reduces dependability

- 85% of Windows crashes are caused by third party code in kernel

- interfaces between extension and host are often poorly documented and understood

- no isolation boundary between code and extension

- extension can access non-public interfaces (reflection)

Single-Process Architecture



- All code and data in single address space

- Rely on language-level memory safety to isolate components

- Dynamic code loading and runtime code generation

- Easy data sharing

- Xerox PARC (Cedar, Smalltalk, etc.) and Lisp Machine model

- Java and .NET model as well

- Runtime is single point of failure


- Shared runtime must also meet all applications' requirements

- Rely on garbage collection to reclaim resources


- Finalizers

- Difficult to construct interactions

Applications Are Still Interdependent




Applications Are Still Interdependent



Runtime



Singularity Sealed Processes



Singularity processes are sealed:

- no dynamic code loading
- run-time code generation
- no code present when process starts execution
- extensions execute in separate addresses
- separate closed environments with well-defined interfaces
- no shared memory
- Process is fundamental unit of failure isolation
- Better security, verification, failure handling, optimization

State Benefit Of Sealed Processes

Program	Whole Code	Reachable Code	% Reduction
Kernel	2.27 MB	1.29 MB	46%
IDE Disk Drive	1.81 MB	455 KB	75%
Web Server	1.73 MB	765 KB	56%
Content Extension	2.34 MB	502 KB	77%

- Reduces process code size by up to 75%
- Fewer code paths => better optimization & error analysis

Need for Lightweight Processes

- Existing processes rely on expensive hardware virtual memory and protection mechanisms
 - VM prevents reference to other processes' pages
 - protection prevents unprivileged code from accessing system resources (e.g. VM page tables)
- Processes are expensive to create and schedule
 - encourages monolithic program development
 - large, undifferentiated applications
 - dynamic code loading
 - threading to allow independent control flow

Software Isolated Processes (SIPs)

- protection and isolation enforced by language safety and kernel API design, not hardware
 - no protection over virtual pages
 - all of a process's objects reside in pages (code, data, address space)
 - language safety means process can't create a virtual reference to other pages
- Global Invariants
 - no process contains a pointer to another process's object space
 - no pointers from exchange heap into process



Interprocess Communications

- Channels are strongly typed (value and behavior)
- bidirectional communications ports
- Messages are in the runtime language support
- Messages live outside processes, in exchange heap
 - only a single message
- "Mailbox" semantics enforced by linear types



Failure Isolation

- SIPs are failure containers
 - no shared implementation or state across SIPs
 - process runtimes are distinct
- On SIP failure
 - clean failure notification on peer channel endpoints
 - resources reclaimed by OS
- Recovery feasible, not automatic or transparent
 - peers can recover and continue

Porting Your System To A Type System

- Process integrity depends on type and memory safety
 - currently trust compiler and runtime
- TAL can eliminate compiler from trusted computing base
- Working on verifying the GC as well

Sing#
C#
source



MSIL
+



Sing#

Singularity
system

compiler
verification

byte code
verification

Our System to A Type System

- Process integrity depends on type and memory safety
 - currently trust compiler and runtime
- TAL can eliminate compiler from trusted computing base
- Working on verifying the GC as well




Hardware Protection Is Orthogonal






Hardware And Software Isolation

Baseline



Workload



Micro Benchmarks


Metric	Singularity	Other CPUs (cycles)		
		Core 2 E6700 Intel 64	Core 2 E6700 Intel 64	Core 2 E6700 Intel 64
Minimum time API call	80	170	170	170
Message request/reply	1,041	1,300	1,300	1,300
Process create & start	388,000	1,030,000	1,030,000	1,030,000

Why?


- all SIPs run in ring 0
- static verification replaces hardware protection
- good optimizing compiler not all

More Verification


- Integrate specifications throughout system development language
- Interprocess communication
- System configuration
- Detect errors early, verify code late
- Language safety essential to system integrity



Flashdisk



Virus Scan



More Verification

Integrate specifications throughout system

- language

- interprocess communication

- system configuration

Detect errors early, verify code late

- language safety essential to system integrity



Channel Contracts

using System;
using System.Net; // TopSocketContract

```
public class TopSocketContract : IChannelContract
{
    public void Read(Socket socket, byte[] buffer)
    {
        // ...
    }

    public void Write(Socket socket, byte[] buffer)
    {
        // ...
    }

    public void Close(Socket socket)
    {
        // ...
    }

    public void GetRemoteClose(Socket socket)
    {
        // ...
    }

    public void GetNoMoreData(Socket socket)
    {
        // ...
    }
}
```

```
try { Read(socket, buffer); }
catch { }


try { Write(socket, buffer); }
catch { }

try { Close(socket); }
catch { }

try { GetRemoteClose(socket); }
catch { }

try { GetNoMoreData(socket); }
catch { }
```

? = receive
! = send



Channel Contracts

Contract

API
Contract

```

// ChannelContract.cs
using System;
using System.Collections.Generic;
using System.Linq;
using System.Text;
using System.Threading.Tasks;

namespace ChannelContracts
{
    public interface IChannelContract
    {
        void SendRead();
    }

    public class ChannelContract : IChannelContract
    {
        public void SendRead()
        {
            // ...
        }
    }
}

```

Channel Contracts

```
it is a TcpConnectContract {  
  state Connected : one {  
    Read? -> @readResult  
    Write? -> @wResult  
  
    GetLocalAddress? -> IPAddress1 :  
      Connected,  
    GetLocalPort? -> Port1 : Connected,  
  
    DoneSending? -> @receiveOnly,  
    DoneReceiving? -> @sendOnly,  
    Close? -> Closed,  
    Abort? -> Closed;  
  }  
  
  state Reading : one {  
    Data? -> Connected;  
    WriteData? -> @sendOnly  
    RemoteClose? -> @zombie  
  }  
  
  conn.SendReadQ,  
  " "  
  " " conn.Data(readData)  
  dataBuffer.Add(1)(readData),  
  " " true,  
  
  use conn.RemoteClose() :  
    return false;  
}
```

Missing Case

case conn.NoMoreData

Contract conformance statically detects
subtle errors such as deadlock

Applications Specifications

- Application is first-class abstraction with identity
 - code
 - resources
 - manifest
- Manifest specifies
 - software components
 - dependencies
 - exported channels
 - hardware or software resource requirements

Device Driver Specification

```
class S3Tri064Config : DriverCategoryDeclaration {
    [IoMemoryRange] framebuffer;
    [IoMemoryRange] textBuffer;
    ...
    [IoPortRange] control;
    TRef<ExtensionContract Exp Start> pnp;
    TRef<ServiceProviderContract Exp Start> video;
    ...
}
```

requires PCI Device

requires 4MB frame buffer
declared in PCI config

requires system
console buffer

requires VGA I/O
ports


requires channel to
parent process for
control

provides channel
for clients to access
video device

Specification Used In Many Ways

Driver
(Source +)

Specification Used In Many Ways




Specification Used In Many Ways



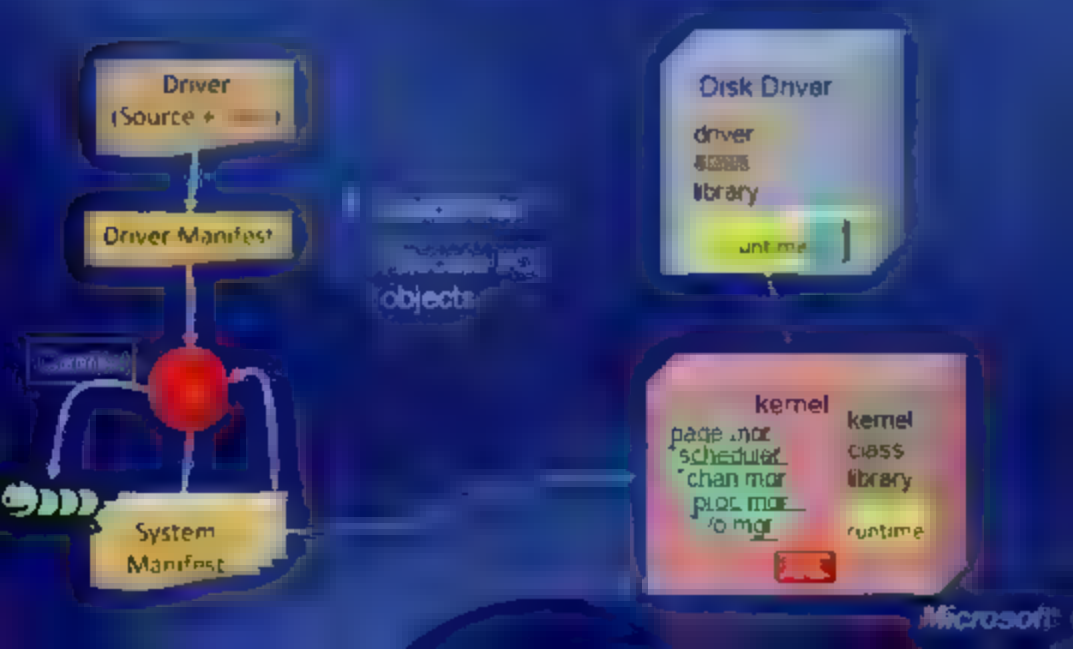
Specification Used In Many Ways



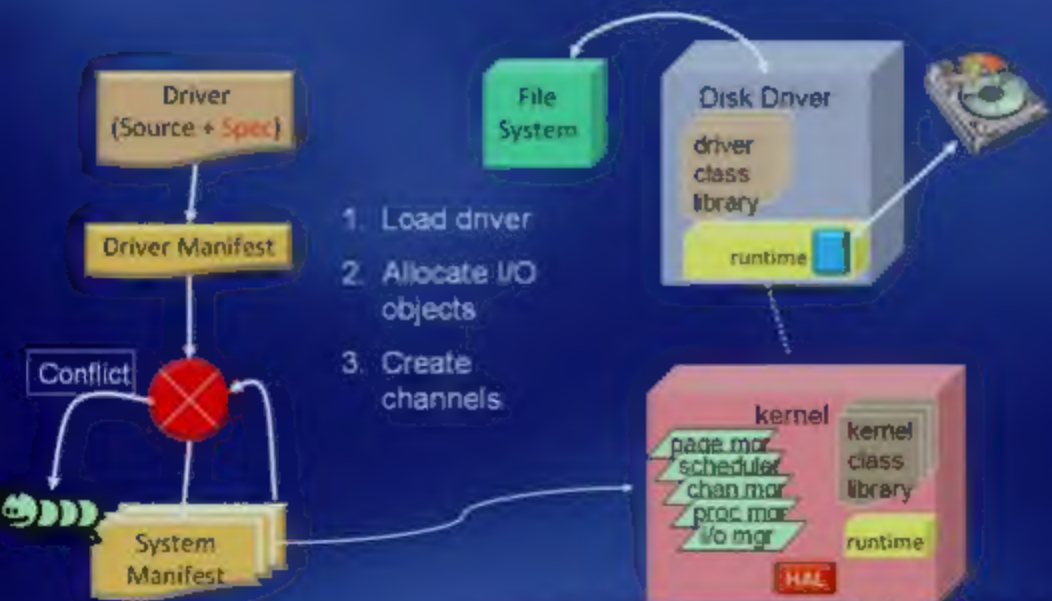
Specification Used In Many Ways



Specification Used In Many Ways



Specification Used In Many Ways



Verification Of System Configuration

- Verification ensures
 - never install an program that will break another program
 - never start a program without appropriate resources
 - never grant a program access to undeclared resources
- All of these checks performed statically

Device Driver Specification

```
[DriverCategory]
[signature("/pci/03/00/5333/8811")]
class S3Trio64Config : DriverCategoryDeclaration
{
    [IoMemoryRange(0, Length = 0x400000)]
    IoMemoryRange framebuffer;

    [IoFixedMemoryRange(Base = 0xb8000, Length = 0x8000)]
    IoMemoryRange textBuffer;

    ...

    [IoFixedPortRange(Base = 0x3c0, Length = 0x20)]
    IoPortRange control;

    [ExtensionEndpoint(typeof(ExtensionContract.Exp))]
    TRef<ExtensionContract.Exp;Start> pnp;

    [ServiceEndpoint(typeof(VideoDeviceContract.Exp))]
    TRef<ServiceProviderContract.Exp;Start> video;

    ...
}
```

requires PCI Device

requires 4MB frame buffer
(declared in PCI config)

requires system
console buffer

requires VGA I/O
ports

requires channel to
parent process for
control

provides channel
for clients to access
video device

Summary

- Singularity is basis for more dependable systems
 - pervasive use of safe programming languages
 - lightweight, closed, customizable run-time environment
 - verifiable specification of system behavior
- Working research prototype
 - driving research in large number of areas
- More information:
 - <http://research.microsoft.com/os/singularity>
 - Growing number of TRs & papers